Customer Learning and Revenue Maximizing Trial Provision

Takeaki Sunada

Department of Economics: University of Pennsylvania

December 27, 2017

Trial products

- Trial (demo) product: common practice in selling experience-durable goods.
  - Information goods: Software, Music, Drama series
  - Subscription service: Amazon Prime, Fitness gym membership
- Wide variety of product design:

  “Time-locked” (Limited-time): Matlab, Microsoft Office
  “Functionality-limited”: Dropbox, Adobe Photoshop

Variations in trial offerings

Variations in trial offerings

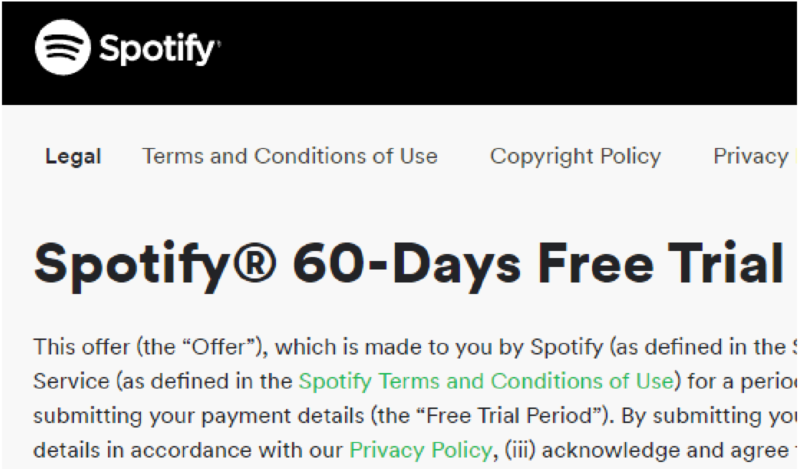
Spotify

Legal | Terms and Conditions of Use | Copyright Policy | Privacy

Spotify® 60-Days Free Trial

This offer (the “Offer”), which is made to you by Spotify (as defined in the
Service (as defined in the Spotify Terms and Conditions of Use) for a perio
submitting your payment details (the “Free Trial Period”). By submitting yo
details in accordance with our Privacy Policy, (iii) acknowledge and agree

Variations in trial offerings

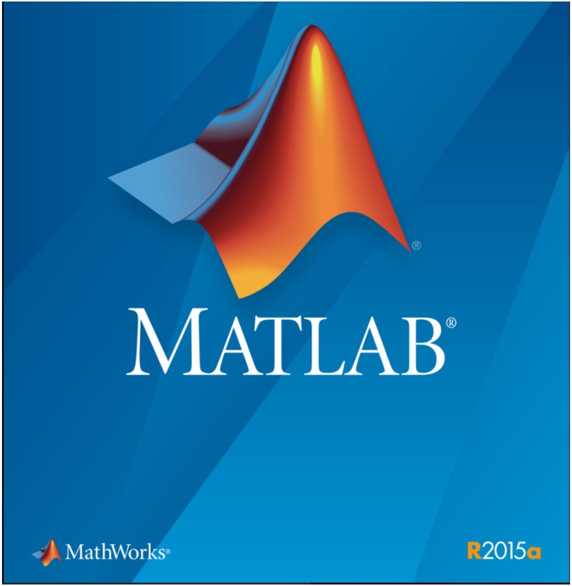

Variations in trial offerings

Variations in trial offerings

Variations in trial offerings

10 hour trial for Madden NFL 18 now available through EA Access on Xbox One

by pastapadre

Posted August 16th, 2017 at 11:15 am

いいね！

Research question

Research question

- What is the optimal trial design (duration of free usage, accessible functionalities) to maximize firm revenue?
- What are the determining factors?

Research question

Research question

- What is the optimal trial design (duration of free usage, accessible functionalities) to maximize firm revenue?
- What are the determining factors?

Why care?

- A firm manager: lack of guideline for trial product design
- A researcher: a particular environment to evaluate the impact of product information provision on the demand

Information provision: what literature has to say

- Literature agrees that trial provision impacts the customer willingness-to-pay through learning-by-using.

but,

- “Whether” and “how” to provide a trial: no unique prediction.
- Optimal trial depends on the nature of customer learning-by-using.
  - Speed of learning, size of initial uncertainty, etc
  - Need for empirical investigation

This paper

1. Build and estimate a model of customer learning.
   - First empirical analysis of customer learning-by-using.
2. Identify the learning mechanism: how trial provision impacts the demand.
3. Counterfactual: find the revenue maximizing trial design.

Market: a gaming software

A major sport game software released annually:

- 3D real-time play in the match
  - requires game-specific play skill.
- 4 game modes are available: “functionality” in this setup
  - Creating a dream team, Simulating a player career, etc
- Largest sales in the category by a large margin: assuming monopolist throughout

Related literature

- Consumer learning - Erdem and Keane (1996), Goettler and Clay (2011), Che, Erdem and Öncü (2015), etc.
- Trial product and product demonstration - Lewis and Sappington (1994), Heiman and Muller (1996), Cheng, Li and Liu (2015), etc.
- Software industry - Lee (2013), Gil and Warzynski (2015), Engelstatter and Ward (2016), etc.

Tentative results

1. Customers are risk aversed and face uncertainty: room for firm intervention.
2. Provision of time-locked trial with 5-7 free sessions can be profitable in some cases ($\sim$ 10% revenue increase).
   - However, in vast majority of cases opportunity cost of lost sales domiantes. No trial is optimal.

Roadmap of the talk

1. Key data moments to support the hypothesis: customer learning
2. Illustration of firm strategy and relevant trade-offs
3. Model formulation: how the mechanism maps into the model.

Data: a gaming software

Session record data provided through Wharton Customer Analytics Initiative (WCAI).

- Randomly sampled 4,956 registered users, no trial experience
- Activation date and history of play (how long and which gamemode) from the date of purchase till the end
- "Session": unit of observation. One session consists of a continuous play of a game mode (may contain multiple matches).

From online price comparison website, history of weekly prices.

Data pattern

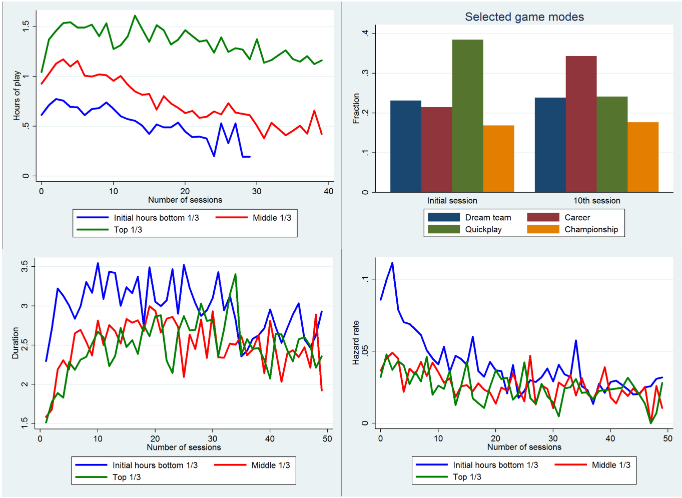

Figure: (Clockwise) Hours per session; Game mode selection; Dropout rate, Duration between sessions

Data pattern

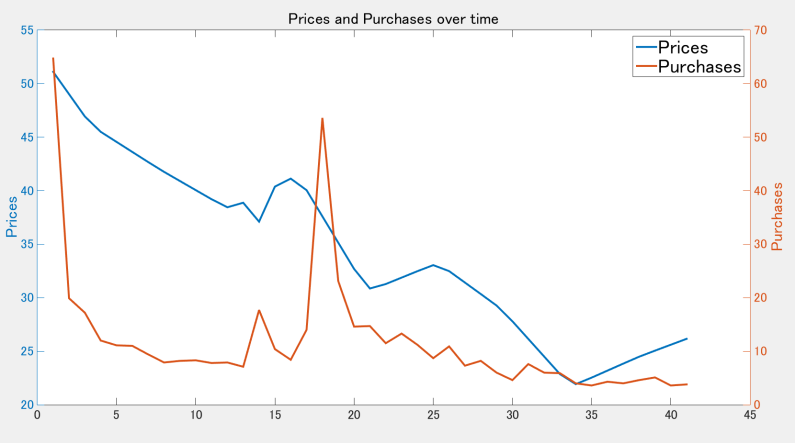

Data pattern

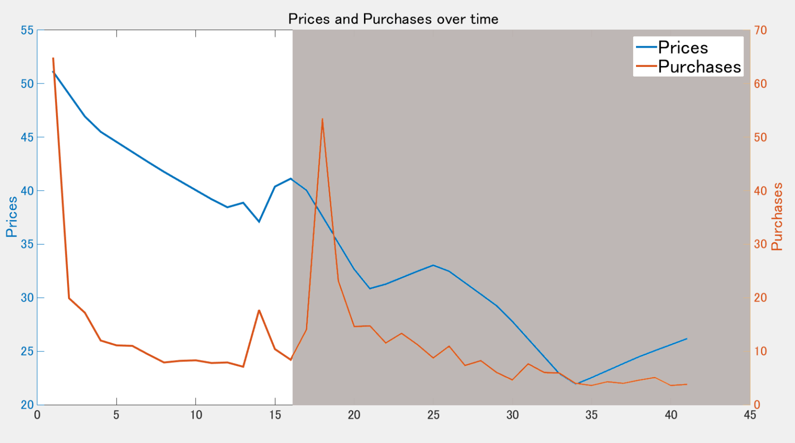

Suggestive evidence for customer learning

In general, customer learning is not separately identified from nonparametric form of preference heterogeneity. But,

1. 42% of users start from "practice mode".
2. 7% initial drop-out rate.

Customer experimenting

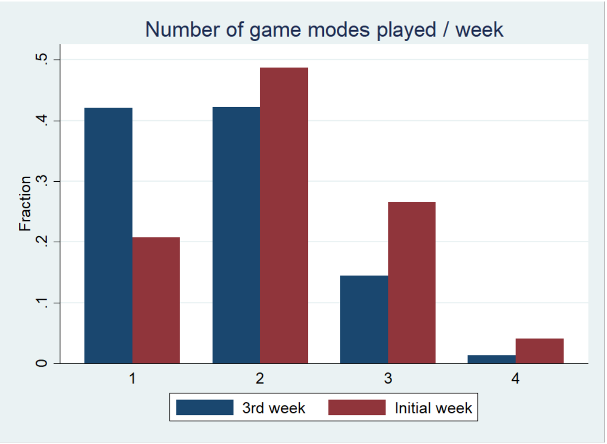

Tendency to play more gamemodes at the beginning - experiment behavior.

Roadmap of the talk

1. Key data moments to support the modeling
2. Illustration of firm strategy and relevant trade-offs
3. Model formulation: how the mechanism maps into the model.

Trial product and the firm objective

Consider customers facing uncertainty about their “match value”.
- Preference, skill, etc

If they are risk aversed, eliminating the uncertainty increases their WTP on average.

The trial design impacts “how” learning occurs: allows firm to manipulate the WTP distribution at purchase.
- Time-locked vs Functionality-limited

Different trade-offs: Time-locked trial

Limiting time essentially limits the number of sessions playable for free.

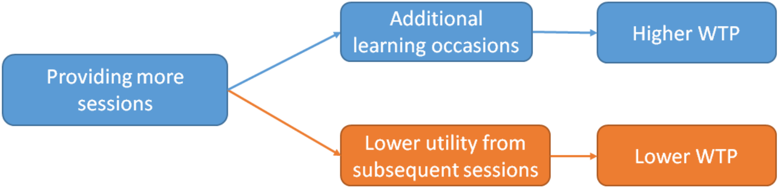

- When initial learning is quick relative to initial utilility decay, an optimal time-limit exists (WTP inverse-U shaped).
- Higher learning speed/utility decay ratio improves the profitability of limited time trial.

Different trade-offs: Limited functionality trial

- Learning spill-over: How much a customer can learn about functionality X when she consumes functionality Y.
- When spill-over is large, providing few functionalities is sufficiently informative about the whole product.
- Otherwise, need for providing many functionalities to facilitate learning - smaller incremental value from the full product.

Different channels

The best practice hence depends on the nature of learning process.

	Large Spillover	Small Spillover
Fast Learning	Limited time and/or Limited functionality	Limited time
Slow Learning	Limited functionality	Not providing trial

The model has to identify these factors, as well as the customer's risk aversion.

Roadmap of the talk

1. Key data moments to support the modeling
2. Illustration of firm strategy and relevant trade-offs
3. Model formulation: how the mechanism maps into the model.

Model

Two dynamic programming: purchase decision, then usage decision

Usage decision: Bayesian learning model with forward-looking customers
- They know they learn. Trade-off between flow payoff vs future return

Purchase decision: forward-looking customers' purchase timing choice.
- The value function from usage decision problem used as an input

The model for usage

- Finite horizon dynamic programming with terminal period T.
- True preference for each functionalities $\theta_i = \{\theta_{i1}, ..., \theta_{im}, ..., \theta_{iM}\} \sim N(\mu, \Sigma)$, unknown to the customer.
- At $t = 0$, She starts from an initial belief $b_{\theta i0} = \{\mu_{i0}, \Sigma_0\}$.

Model timeline

- At each session t, a user chooses the functionality *m* and the hours of usage
  - Sequential order: functionality (dynamic), and then hours (static).
- After the session, receive an an informative signal for the chosen functionality, update the belief $b_{\theta it}$.
- Two random variables determine whether she stays active or terminates, and the duration till next session if active.
  - Interpretable as a reduced form policy.

Choice of the hours

Static expected utility maximization;

$$\max_{x_{imt}} E(u(x_{imt}, \theta_{im}, \nu_{imt}, h_t)|b_{\theta it})$$
$$= f(b_{\theta it})x_{imt} - \frac{(\gamma_1 \nu_{imt} + \gamma_2 \nu_{imt}^2 + x_{imt})^2}{2(1 + \alpha_1 h_t)},$$

$$\text{where } f(b_{\theta it}) = (E(\theta_{im}^{\rho}|\theta_{im} > 0, b_{\theta it})P(\theta_{im} > 0|b_{\theta it}))^{\frac{1}{\rho}}, \ \rho > 0.$$

x_{imt}: the hours of usage of functionality m at session t
ν_{imt}: the number that functionality m was chosen in the past t-1 sessions
h_t: holiday indicator

- If risk aversed, then smaller variance $\rightarrow$ higher utility, longer play hours.

Optimal hours of play and game mode specific utility

$$v_{imt}(b_{\theta it}, \nu_{imt}, h_t)$$
$$= \max\left\{\frac{f(b_{\theta it})^2(1+\alpha_1 h_t)}{2} - f(b_{\theta it})(\gamma_1\nu_{imt} + \gamma_2\nu_{imt}^2), 0\right\}.$$

$$x^*_{imt}(b_{\theta it}, \nu_{imt}, h_t) = \max\left\{f(b_{\theta it})(1+\alpha_1 h_t) - (\gamma_1\nu_{imt} + \gamma_2\nu_{imt}^2), 0\right\}.$$

ν_{imt}: the number that functionality m was chosen in the past t-1 sessions
h_t: holiday indicator

The choice of functionality

Let $\Omega_{it} = (\{b_{\theta it}\}_{m=1}^{M}, \{\nu_{imt}\}_{m=1}^{M}, h_t)$ be the state.

$$\begin{aligned} V_{it}(\Omega_{it}) &= E(\max_{m_{it}} v_{imt}(b_{\theta it}, \nu_{imt}, h_t) + \epsilon_{imt}\sigma_{\epsilon 1} \\ &\quad + E(\beta(\Omega_{i,t+1}) V_{i,t+1}(\Omega_{i,t+1}) | \Omega_{it}, m_{it})) \\ &= \sigma_{\epsilon 1} \log \left(\sum_m \exp \left(\frac{1}{\sigma_{\epsilon 1}} (v_{imt}(b_{\theta it}, \nu_{imt}, h_t) \right.\right. \\ &\quad \left.\left. + E(\beta(\Omega_{i,t+1}) V_{i,t+1}(\Omega_{i,t+1}) | \Omega_{it}, m_{it})) \right) \right). \end{aligned}$$

Discount factor $\beta(\Omega_{i,t+1})$: including the probability of drop-out and future frequency of play.

v_{imt} already level-scale normalized: no need to normalize $\sigma_{\epsilon 1}$.

Belief update

Initial belief: $b_{\theta i0} = \{\mu_{i0}, \Sigma_0\}$.
After each session t, she receives an unbiased signal $s_{imt}|\theta_{im} \sim N(\theta_{im}, \sigma_s^2)$ for the chosen m.

$$\mu_{i,t+1} = \mu_{it} + \Sigma_{it} Z_{it}' (Z_{it} \Sigma_{it} Z_{it}' + \sigma_s^2 * I)^{-1} * (s_{imt} - \mu_{imt}),$$
$$\Sigma_{i,t+1} = \Sigma_{it} - \Sigma_{it} Z_{it}' (Z_{it} \Sigma_{it} Z_{it}' + \sigma_s^2 * I)^{-1} Z_{it} \Sigma_{it},$$

where Z_{it} represents the functionality chosen at t by i. Learning spill-over exists: preference correlation.

The model for purchase decision

- Forward looking consumer with rational expectation for tomorrow's price: an optimal stopping problem
- Buy when the value from buying $>$ the value from waiting.
  - Value from buying: $V(\Omega_0)$ (value function derived from the model of usage)
  - Value from waiting: future price reduction

$$V_p(\Omega_{i0}, p_\tau) = E(\max\{V(\Omega_{i0}) - \alpha_p p_\tau + \epsilon_{1\tau}\sigma_{\epsilon 2}, \\ \beta E(V_p(\Omega_{i0}, p_{\tau+1})) + \epsilon_{0\tau}\sigma_{\epsilon 2}\}),$$

where τ is a calendar week.

The model formulation and the research objective

Model and trial product provision:

- If risk aversed, value function *increases* after a few sessions.
  - Rationale for providing trials
- Two dimensions of learning: over time and across functionalities. Limiting either one impacts learning.
- Firm trade-off easily attributable to model parameters:
  - Speed of learning: signal precision (σ_s)
  - Boredom: utility decay (γ) and dropout probability
  - Learning spill-over: preference correlation (Σ)

Simulated maximum likelihood:

- Likelihood of usage = probability of game mode choice $\times$ hours of play $\times$ duration between sessions $\times$ termination probability. Multiply over sessions.
- Likelihood of purchase = probability of observing the purchase pattern in the data.
  Customer arrival process: A peak at the week of product launch, then uniform arrival afterward.
- Combine both likelihood and integrate over unobservables.

Model fit - usage

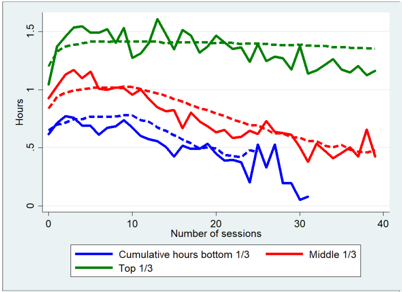

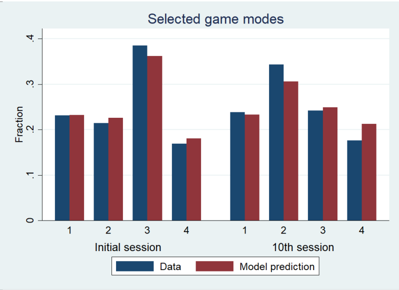

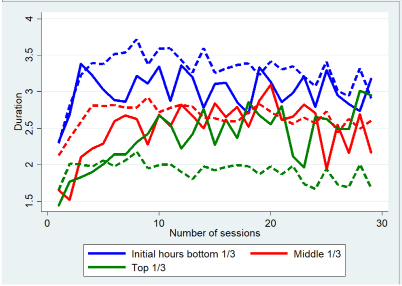

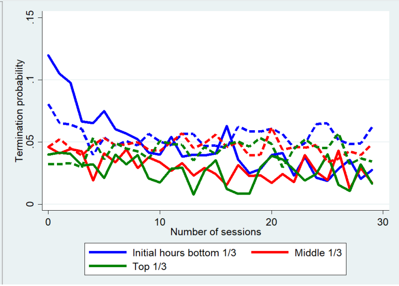

Model fit - purchase

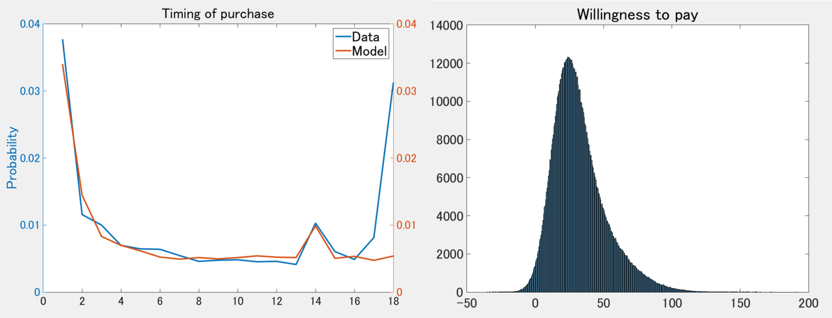

Summary of findings

- Customers are risk aversed: $\rho = 0.5802 < 1$.
- Large initial uncertainty: 95% confidence interval of someone with \$40 initial mean WTP $= [30.8271, 49.1729]$.
- Learning is quick:
  "s.d. of signal : s.d. of initial belief = 1 : 4".

Counterfactual - time-locked trial

- Suppose that $\tilde{T} < T$ free sessions is provided, along with the full product.
- At period 0, customers can pick either to take trial or to buy the full product.
- While using a trial, they can decide on whether to make a purchase or not at the end of each calendar day.
- Value functions for play and purchase are solved jointly.
- The solution of the dynamic programming yields an aggregated demand function.
- The firm maximizes the revenue (price$\times$aggregated demand, summed over periods) by choosing $\tilde{T}$.

Provision of time-locked trial is similar, except $\tilde{M} < M$ rather than $\tilde{T} < T$.

Counterfactual - time-locked trial

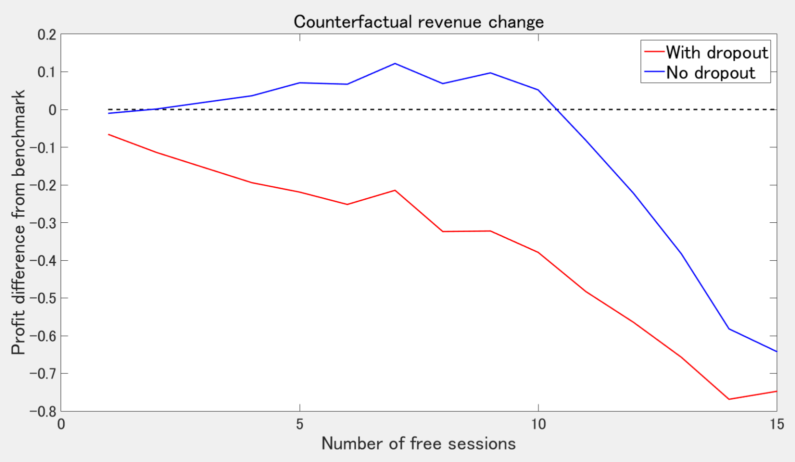

WTP of those who survive does increase. However, high initial drop-out rate offsets all the gains.

Conclusion

- This paper has addressed a question of optimal trial provision.
- Customers are indeed risk-aversed and face uncertainty at purchase. Provision of trial increases WTP.
- However, initial high drop-out rate makes the provision of time-locked trial suboptimal in majority of the parameter range.